

MULTIPROTOCOL LABEL SWITCHING (MPLS) EDGE SERVICE EXTRACTION

5

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application is entitled to the benefit of provisional Patent Application
10 Serial Number 60/398,001, filed 22 July 2002.

FIELD OF THE INVENTION

15 [0002] The invention relates generally to a technique for managing network traffic,
and more particularly, to a technique for providing network services to customers through
a provider network that includes, for example, a multiprotocol label switching (MPLS)
domain.

20

BACKGROUND OF THE INVENTION

[0003] Virtual private LAN services (VPLS) over MPLS provides LAN-like
connectivity between geographically diverse customer locations. Draft standards for
25 implementing VPLS over MPLS are presented in “Virtual Private LAN Services over
MPLS”, IETF draft-lasserre-vkompella-ppvnp-vpls-04.txt, March 2003, and “Virtual
Private LAN Service”, IETF draft-kompella-ppvnp-vpls-02.txt, May 2003, “Transport of
Layer 2 Frames over MPLS”, IETF draft-martini-l2circuit-trans-mpls-09.txt, April 2002,
and “Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS
30 networks”, IETF draft-martini-l2circuit-encap-mpls-04.txt, November 2001, all of which
are incorporated by reference herein.

[0004] The basic operation of VPLS is described with reference to Fig. 1 and involves establishing virtual circuit (VC) label switching paths (LSPs) and tunnel LSPs. Fig. 1 depicts an MPLS domain 102, two service provider edge devices (PEs) 104, geographically diverse customer locations 106 for two different customers, customer A and customer B, and the corresponding tunnel LSP and VC LSPs (e.g., VC_A LSP and VC_B LSP). In operation, a customer's Ethernet packet is either switched or routed by a customer device to one of the PEs (also known as MPLS label edge router (LERs)). The respective PE classifies the packet based on either the incoming port or the virtual local area network (VLAN) identified (ID) of an IEEE 802.1q tagged packet. The packet is then mapped to a user-defined Forwarding Equivalence Class (FEC), which specifies how the packet gets forwarded. The FEC lookup yields the outgoing port of the packet and labels that are used to encapsulate the packet. Fig. 2 depicts an example of a frame encapsulation format for implementing VPLS. A complete description of the frame format is described in the above referenced IETF document entitled "Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS networks". The frame format includes the original Ethernet packet 208, an MPLS label stack 209, and an outer Ethernet header 212. The label at the top of the MPLS label stack is the tunnel label, which is used to transport the packet across the provider's MPLS domain. The label at the bottom of the MPLS label stack is the VC label, which is used by the egress PE to determine how to process the packet.

[0005] Upon transport through the MPLS domain, backbone label switch routers (LSRs) (not shown) in the MPLS domain only look at the tunnel label to switch the labeled packet through the MPLS domain. It is possible that additional labels get pushed along the way. The tunnel label at the top of the MPLS label stack is removed at the penultimate hop (i.e., the hop prior to the egress PE) and the packet is passed to the egress PE with only the VC label. The egress PE uses the VC label to determine how to process the packet. The packet is then forwarded to the outgoing port that is identified via the VC label.

[0006] Fig. 3 is an expanded view of a PE 204 that logically depicts an example of how customer-specific VPLS traffic is passed through the PE into an MPLS domain. In the example, customers A and B are connected to the PE via customer-specific ports

(e.g., ports 1, 2, and 3). VLAN traffic received on the customer-specific ports is associated with a customer-specific VC LSP and a tunnel LSP. In the example, each customer has a unique VC LSP (e.g., VC_A LSP and VC_B LSP) and traffic that is destined for the same PE (e.g., PE₂) uses the same tunnel LSP.

5 [0007] In many applications, VPLS involves more than a point-to-point connection between two PEs. As depicted in Fig. 4, VPLS often involves connections between multiple distinct customer locations 406. For example, customers A and B may each have three distinct locations that are connected via the MPLS domain at PEs 1, 2, and 3. The three PEs are connected by three bidirectional tunnel LSPs (where each bidirectional
10 tunnel is formed by two unidirectional LSPs). Customer-specific VC LSPs, which connect each PE, are also established for each customer between the connected PEs. For example, at PE₁, for customer A, there is a VC LSP that connects PE₁ to PE₂ and a different VC LSP that connects PE₁ to PE₃. For customer A, one VC LSP is used to carry traffic from PE₁ to PE₂ and the other VC LSP is used to carry traffic from PE₂ to PE₃.
15 The full mesh of customer-specific VC LSPs enables unique broadcast domains for customers A and B. Although the customers use customer-specific VC LSPs, the customers share the same tunnel LSPs for transport between the same PEs.

[0008] While the above-described VPLS scheme works well to provide “virtual private” or “transparent” LAN services, customers often demand more than LAN services
20 from their service provider. For example, customers often need access to other network services such as Internet, video on demand (VoD), and PSTN services. Current VPLS technologies do not adequately address the delivery of other network services. Therefore, what is needed is a technique for providing customers with services, such as Internet, VoD, and PSTN services, in addition to “virtual private” or “transparent” LAN services
25 that is flexible and efficient to implement and that is complementary with the emerging VPLS standards.

SUMMARY OF THE INVENTION

[0009] In accordance with the invention, a set of VLAN IDs is explicitly identified for use with a first service. The rest of the customer traffic is considered as part of a default service. Traffic that is received at a PE from a customer is examined to identify whether or not the traffic belongs to the first service. For example, the VLAN ID and incoming port of a packet is compared to the set of VLAN IDs that were allocated to the first service on the respective port. Traffic that is identified as belonging to the first service is “extracted” from the default service and forwarded on a path that is related to the first service. The remaining traffic is forwarded on a path that is related to the default service. In an embodiment, the service extraction technique is implemented across an MPLS domain that utilizes the VPLS over MPLS techniques that are described in the above-identified IETF draft standards. An advantage of the service extraction technique is that specific services can be separated from a default service by explicitly defining the VLAN IDs that are included in the service without having to redefine the default traffic class. Extracting desired services without having to redefine the default traffic class enables flexible and efficient network management.

[0010] In an embodiment, the default service is a virtual private LAN (VPL) service and the first service is a non-VPL service such as Internet, VoD, or PSTN. In this embodiment, the set of VLAN IDs is explicitly allocated for use with the non-VPL traffic. The rest of the customer traffic is considered VPL traffic. Traffic that is received at a PE from a customer is examined to identify the non-VPL traffic. For example, the VLAN ID of a packet is compared to the set of VLAN IDs that were allocated to non-VPL traffic. Traffic that is identified as non-VPL traffic is extracted from the VPL traffic and sent to a service gateway. The remaining traffic is forwarded within the customer’s VPL.

[0011] Other aspects and advantages of the present invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings, illustrating by way of example the principles of the invention.

BREIF DESCRIPTION OF THE DRAWINGS

- 5 **[0012]** Fig. 1 depicts an MPLS domain, two service provider edge devices (PEs), and geographically diverse customer locations for two different customers, customer A and customer B.
- [0013]** Fig. 2 depicts an example of a frame encapsulation format for implementing VPLS.
- [0014]** Fig. 3 is an expanded view of a PE that logically depicts how customer-specific VPLS traffic is passed through the PE into an MPLS domain.
- 10 **[0015]** Fig. 4 depicts an MPLS domain and three PEs that are used to provide VPLS to two different customers.
- [0016]** Fig. 5 depicts a service provider MPLS domain that connects multiple geographically diverse customer locations to each other and to an additional network, such as the Internet.
- 15 **[0017]** Fig. 6 is an expanded view of a PE that logically depicts non-VPL traffic that is extracted from VPL traffic at a PE upon insertion into the MPLS domain in accordance with an embodiment of the invention.
- [0018]** Fig. 7 graphically depicts an example of the service extraction process in accordance with an embodiment of the invention.
- 20 **[0019]** Fig. 8 is an expanded view of a PE that logically depicts VPL and extracted non-VPL traffic that is sent out of the same physical port but different tunnel LSPs in accordance with an embodiment of the invention.
- [0020]** Fig. 9 depicts a network architecture in which a customer is provided access to distinct VoD and Internet networks in addition to the VPL services in accordance with an embodiment of the invention.
- 25 **[0021]** Fig. 10 depicts an example of two different customers that access the same non-VPL service in accordance with an embodiment of the invention.
- [0022]** Fig. 11 depicts an embodiment of a method for managing network traffic in accordance with an embodiment of the invention.
- 30 **[0023]** Fig. 12 depicts another embodiment of a method for managing network traffic in accordance with an embodiment of the invention.

[0024] Fig. 13 depicts a system for implementing service extraction within a PE in accordance with an embodiment of the invention.

[0025] Throughout the description, similar reference numbers may be used to identify similar elements.

5

DETAILED DESCRIPTION OF THE INVENTION

[0026] Fig. 5 depicts a service provider MPLS domain 502 that connects multiple geographically diverse customer locations 506 to each other and to an additional network 514, such as the Internet via PEs 504. In accordance with the invention, the customers (e.g., customers A and B) have access to virtual private LAN (VPL) and Internet services through the PEs and the MPLS domain. Providing VPL services and additional network services involves establishing different tunnel LSPs between the PEs for the VPL services and the additional services as is described in the above-identified IETF draft standards. Throughout the description, network services and traffic that are not part of the VPL services are referred to as “non-VPL” services or “non-VPL” traffic. Non-VPL services may include, for example, access to the Internet, a video on demand (VoD) network, a publicly switched telephone network (PSTN), or some other specialty network. Referring to the example of Fig. 5, the non-VPL services are accessed through PE₄ (also referred to herein as a “service gateway”). Non-VPL tunnels are established between PE₄ and the customer-connected PEs (PE₁, PE₂, and PE₃) to provide access to the Internet. VPL services are provided via VPL tunnels that are established between the customer-connected PEs (PE₁, PE₂, and PE₃). For description purposes, non-VPL and VPL tunnels are bidirectional tunnels that are actually formed by two unidirectional MPLS LSPs as is described in the above-referenced IETF draft standards.

[0027] In accordance with an embodiment of the invention, a set of VLAN IDs is explicitly identified for use with non-VPL traffic. The rest of the customer traffic falls into a default category and is treated as VPL traffic. Traffic that is received at a PE from a customer is classified to determine if the traffic is non-VPL traffic. For example, the VLAN ID of a packet is compared to the set of VLAN IDs that were explicitly allocated to non-VPL traffic. In an embodiment, the VLAN ID comparison is made on a per-port

or per-customer basis such that all of the VLAN IDs have per-port or per-customer significance. Traffic that is identified as non-VPL traffic is “extracted” from the default service (e.g., the VPL service) and sent to a service gateway. The process of extracting traffic is described in more detail below. The remaining traffic is forwarded within the customer’s VPL. In an embodiment, the incoming traffic goes through a traffic classification process, which identifies the traffic as belonging to a non-VPL or a VPL traffic class. The traffic classification process may be VLAN based or VLAN and port (or customer) based. Additionally, the classification could be based on a VLAN range, or ranges, or a port/VLAN range combination. Classifying the incoming traffic based on port and VLAN gives VLAN IDs per-port significance and enables the full range of VLANs to be available on each port (or to each customer).

[0028] Fig. 6 is an expanded view of a PE 606 that logically depicts an example of non-VPL traffic that is extracted from VPL traffic at the point where the traffic is inserted into the MPLS domain. In the example, customers A and B are connected to the PE via customer-specific ports (ports 1 and 2). With regard to customer A, traffic is sent to the PE using VLANs 100, 200, 210, and 220. For example purposes, it is assumed that VLAN 100 is explicitly designated for non-VPL traffic. In operation, the VLAN ID of each packet that is received at the PE from the customer-facing port (e.g., a port that is directly connected to a customer device such as a switch or router) is examined. Packets with the non-VPL VLAN ID (e.g., VLAN 100) are extracted from the VPL traffic. The extracted non-VPL traffic is associated with a virtual channel (VC) LSP that is different from the VC LSPs that are associated with the VPL traffic. The extracted non-VPL traffic is also associated with a tunnel LSP that connects the PE with the desired service gateway. For example, the non-VPL traffic for customer A that is depicted in Fig. 6 is associated with a customer-specific VC LSP 620 and a tunnel LSP 624 that connects the traffic to the service gateway. Specifically, customer A’s VLAN 100 traffic is associated with a customer-specific VC LSP and a tunnel LSP that connects the traffic to the service gateway (e.g., PE₄ as depicted in Fig. 5). Within each packet that is encapsulated as described with reference to Fig. 2, the customer-specific VC LSP and tunnel LSP are identified by corresponding VC and tunnel labels that are included in the MPLS label stack of the packet.

[0029] The traffic that is not explicitly extracted from the VPL is implicitly associated with a customer-specific VC LSP and a tunnel LSP that connect the traffic to the default service (e.g., the VPL service). In the example of Fig. 6, the default traffic is associated with a customer-specific VC LSP and a tunnel LSP that connects the traffic to one of the PEs in the VPL (e.g., PE₂ or PE₃ from Fig. 5). Specifically, customer A's VLAN 200, 210, and 220 traffic is associated with a customer-specific VC LSP 630 and a tunnel LSP 634 that connects the traffic to, for example, PE₂ (as depicted in Fig. 5). The customer-specific VC LSP and tunnel LSP are identified by corresponding VC and tunnel labels that are included in the MPLS label stack. It should be noted that the entire default service can be identified and tunneled using a single VC ID to VC label mapping. Managing all of the default traffic using a single VC ID greatly reduces the amount of VC ID signaling that is required between PE devices. Additionally, it conserves the consumption of available VC IDs.

[0030] With regard to customer B, traffic is sent to the PE using VLANs 150, 210, 220, and 230. For example, it is assumed that VLAN 150 is explicitly designated for non-VPL traffic. At the PE, customer B's VLAN 150 traffic is explicitly associated with a customer-specific VC LSP 622 and a tunnel LSP 624 that connects the traffic to the service gateway (e.g., PE₄ as depicted in Fig. 5). Customer B's VLAN 210, 220, and 230 traffic is implicitly associated with a customer-specific VC LSP 632 and a tunnel LSP 634 that connects the traffic to the default service, for example, PE₂ (as depicted in Fig. 5).

[0031] Note that different VC LSPs are used for a customer's VPL and non-VPL traffic (e.g., VC LSP 620 and VC LSP 630 for customer A). However, the same tunnel LSP (e.g., tunnel LSP 624) is used to transport the non-VPL traffic for customers A and B to PE₄. Likewise, the same tunnel LSP is used to transport the VPL traffic for customers A and B to PE₂. The use of common tunnel LSPs enables efficient signaling and scaling of the system. Also note that the combination of classifying traffic on an incoming port/VLAN ID basis and the VPLS tunneling scheme enables each customer to utilize the total available VLAN ID space (e.g., 4,096 unique 802.1q VLAN IDs). For example, customers A and B are allowed to use the same VLAN IDs (e.g., VLANs 210 and 220) without having their traffic mixed up.

[0032] In accordance with an embodiment of the invention, non-VPL traffic is extracted from VPL traffic using a default technique. According to an embodiment of the default technique, the entire range of VLAN IDs is made available to a customer for use at a customer-specific port. A set of the available VLAN IDs is then explicitly identified for use with non-VPL traffic. The set of VLAN IDs could include only one VLAN ID or multiple contiguous (e.g., VLANs 1 – 7) or non-contiguous (e.g., VLANs 1, 3, and 7) VLAN IDs. Preferably, the set of VLAN IDs is explicitly allocated on a per-port or per-customer basis. In an embodiment, the VLAN IDs are the VLAN IDs that are defined in the IEEE 802.1q specification. Traffic that is identified by the explicitly identified VLAN IDs is considered to be non-VPL traffic. In an embodiment, explicitly identifying VLAN IDs involves coding the identified VLAN IDs into the classification process on a per-port basis as VLAN IDs that are to be handled according to a set of rules that is different from the default traffic. In an embodiment, explicitly identifying VLAN IDs triggers the establishment of corresponding VC IDs, layer 2 (L2) FECs, tunnel labels, and VC labels, where layers are defined by the International Standardization Organization (ISO) in the Open System Interconnection (OSI) model. The explicit allocation of VLAN IDs also triggers a signaling process to other PEs in the MPLS domain.

[0033] Any other traffic (either 802.1q tagged or untagged traffic) received at the port is implicitly considered to be VPL traffic. That is, traffic that is not explicitly identified as traffic that is to be extracted from the VPL traffic falls into the default category and is forwarded within the default service.

[0034] According to the default technique, incoming traffic is first checked on a per-port basis to see if it corresponds to a non-VPL VLAN ID. Traffic having a non-VPL VLAN ID is associated with a corresponding L2 FEC. All of the traffic that is not explicitly identified is implicitly associated with a different corresponding L2 FEC. Fig. 7 graphically depicts an example of the extraction process. In the example of Fig. 7, it is assumed that the entire range of 802.1q VLAN IDs is available to a customer at a customer-specific port of the PE. The set of VLAN IDs, including 802.1q VLANs 1 – 7 and 4001, is explicitly allocated to non-VPL traffic. Note that the number of unique VLAN IDs available according to the 802.1q standard is actually limited to 4,094 because the value of all ones (0xFFF or 4,095) is reserved and the value of all zeros

(0x000 or 0) indicates a priority tag. All other traffic received at the port, either 802.1q tagged or untagged traffic, falls into the default category and is implicitly considered VPL traffic. In the example of Fig. 7, VPL traffic includes 802.1q VLANs 8 – 4000 and 4002 – 4094 and all untagged (or non-802.1q) traffic that is received at a respective port.

5 [0035] In operation, traffic is associated with a particular L2 FEC in response to a traffic classification process that first checks incoming traffic for non-VPL VLAN IDs. If incoming traffic has one of the explicitly identified VLAN IDs, then the traffic is associated with a corresponding L2 FEC. The corresponding L2 FEC is then used to identify a VC ID, which is in turn used to identify the corresponding VC label and tunnel
10 label pair. As depicted in Fig. 7, the explicitly identified VLAN IDs can correspond to different L2 FECs. Traffic that is not explicitly identified by VLAN ID falls into the default category. The default traffic is associated with a corresponding L2 FEC and the L2 FEC is used to identify a VC ID, which is in turn used to identify the corresponding VC label and tunnel label pair. In accordance with the invention, all of the default traffic
15 can be represented by a single VC ID.

[0036] Although in Figs. 5 – 7 the explicitly defined traffic is described as non-VPL traffic and the default traffic is described as VPL traffic, the VPL/non-VPL distinction is not critical to the invention. That is, the explicitly defined traffic could be related to a first service, which could include any type of service, while the default traffic
20 is related to a second service, which also could include any type of service.

[0037] In the example of Fig. 6, the VPL and non-VPL traffic is sent out of the PE on different physical ports. In some network configurations, the service gateway may be reachable through the same physical port from which other PEs within the VPL are reachable. Non-VPL traffic can be extracted from VPL traffic even though the extracted
25 traffic is sent out of the same physical port of the respective PE. For example, VPL and non-VPL traffic can be sent out of the same physical port using different tunnel LSPs that connect the different traffic to the different PEs. Fig. 8 is an expanded view of a PE 806 that logically depicts VPL and extracted non-VPL traffic that is sent out of the same physical port (e.g., port 3) using different tunnel LSPs. As depicted in Fig. 8, non-VPL
30 traffic from customers A and B that is destined to PE₄ is sent from port 3 of PE₁ using one tunnel LSP (tunnel LSP 624) while VPL traffic from customers A and B that is

destined to PE₂ is sent from the same physical port using a different tunnel LSP (tunnel LSP 634).

[0038] In accordance with an embodiment of the invention, more than one type of service can be provided using the above described technique. For example, a customer
5 may desire access to two different networks in addition to the VPL services. Fig. 9 depicts a network architecture in which a customer is provided access to distinct VoD and Internet networks in addition to the VPL services. In the example of Fig. 9, non-VPL traffic for the VoD network 915 is explicitly allocated VLANs 0 – 7 and traffic for the Internet 914 is explicitly allocated VLAN 4001. The remaining traffic falls into the
10 default (DFT) category and is implicitly treated as VPL traffic. In the example of Fig. 9, non-VPL tunnels (which correspond to VLANs 1 – 7 and 4001) are established between each PE and each of the service gateways for non-VPL traffic and VPL tunnels (which correspond to the default traffic) are established between the PEs (PE₁, PE₂, and PE₃) for the VPL traffic.

[0039] It is also likely that more than one customer will desire access to the same non-VPL network (e.g., the Internet). Fig. 10 depicts an example of two different customers that access the same non-VPL service. In the example, customer A has allocated VLAN 4001 to non-VPL traffic while customer B has allocated VLAN 2 to non-VPL traffic. Non-VPL traffic that is received at a PE from a customer-facing port is
20 transported to the service gateway (i.e., PE₄) using a non-VPL tunnel that connects the respective PE to the service gateway. VPL traffic is sent to the appropriate PE(s) using a VPL tunnel that connects the respective PEs in the VPL. Because incoming traffic can be classified on a port/VLAN ID basis and because traffic is tunneled in customer-specific VC LSPs as described above, each customer-specific port can utilize the entire available
25 set of VLAN IDs without causing traffic to be mixed. Therefore, customers can use the same VLANs to identify the same services without causing different customers' traffic to be mixed. For example, customers A and B can both use the same VLAN to identify non-VPL traffic that is to be delivered to the same service gateway with the service gateway translating the common VLAN ID to two unique VLAN IDs.

[0040] Fig. 11 depicts an embodiment of a method for managing network traffic in accordance with an embodiment of the invention. At step 1102, a set of VLAN IDs is

explicitly identified for use with a first service. At step 1104, traffic from a customer at a PE is received, wherein the PE connects to other PEs via a tunnel-capable network. At step 1106, the received traffic is classified. At step 1108, the traffic is associated with the first service in response to the classification if the traffic has a VLAN ID from the explicitly identified set of VLAN IDs. At step 1110, the traffic is associated with a default service in response to the classification if the traffic does not have a VLAN ID from the explicitly identified set of VLAN IDs.

[0041] Fig. 12 depicts another embodiment of a method for managing network traffic in accordance with an embodiment of the invention. At step 1202, a customer-specific VPL is established through an MPLS domain. At step 1204, a set of VLAN IDs is explicitly identified for use with non-VPL traffic. At step 1206, traffic is received from a customer at a PE, wherein the PE connects to other PEs via the MPLS domain. At step 1208, the received traffic is examined to identify non-VPL traffic. At step 1210, non-VPL traffic is forwarded outside of the customer-specific VPL. At step 1212, the remaining traffic is forwarded within the customer-specific VPL.

[0042] Fig. 13 depicts a system for implementing service extraction within a PE as described above with regard to Figs. 5 – 12. The system includes a receive module 1342, a classification engine 1344, an extraction engine 1346, a VPLS engine 1348, a user interface engine 1350, and a memory 1352 for storing explicitly identified traffic. The receive module performs standard functions for receiving network traffic into the PE, either from customer-facing ports or MPLS ports. The classification engine performs classification functions on incoming traffic. As described above, the classification engine can classify traffic at least in part based on VLAN ID, incoming port, or the combination of VLAN ID and incoming port. The user interface engine provides the user interface that enables VLAN IDs to be explicitly provisioned for extraction. Explicitly identified VLAN IDs can be stored in the explicitly identified traffic memory. The VPLS engine performs VPLS functions as described in the above-identified IETF draft standards. The extraction engine performs the service extraction functions described above with regard to Figs. 5 – 12. In an embodiment, the system elements are embodied in an application specific integrated circuit (ASIC), multiple ASICs, a field programmable gate array (FPGA), or multiple FPGAs. Alternatively, the system elements could be embodied in a

multifunction processor and external memory, or any combination of ASICs, multifunction processors, and memory.

5 [0043] In an embodiment, service extraction can be implemented in a network environment other than MPLS. For example, other transport tunnels that can be used by pseudo-wires (e.g., GRE, L2TP, IPSEC, etc.) can also be used, as long as the originating PE can be identified. That is, service extraction can be implemented in other tunnel-capable networks that utilize, for example, GRE, L2TP, or IPSEC.

10 [0044] Although the traffic is described herein as either VPL or non-VPL traffic, it should be noted that, in an embodiment, the traffic is distinguished in that it is related to different services. For example, the service extraction technique can be utilized to deliver the explicitly identified traffic to a first service (e.g., VoD) and the implicitly identified traffic to a second service (e.g., Internet).

15 [0045] Although specific embodiments of the invention have been described and illustrated, the invention is not to be limited to the specific forms or arrangements of parts as described and illustrated herein. The invention is limited only by the claims.